Semi-Sameness and Consciousness: How Minds Construct Usable Identity in a Non-Repeating Universe

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

December 31, 2025

If no spacetime event ever repeats, then “sameness” cannot be a primitive feature of the world. Yet human life depends on sameness: we re-identify objects, recognize people, keep promises, assign responsibility, and narrate continuity across years. This paper proposes that consciousness is the mechanism that manufactures a workable middle layer between unrepeatable reality and practical agency: artificial semi-sameness. Semi-sameness is not strict identity and not mere illusion. It is a tolerance-scoped equivalence that allows finite agents to compress an unrepeatable event stream into stable handles such as “the same person,” “the same place,” and “the same self.” On this view, consciousness is not only a theater of experience; it is an equivalence-maker that extracts invariants, assigns tolerances according to stakes, binds features into objects and agents, and sustains continuity chains that ground accountability and repair. The paper formalizes this idea without mysticism, distinguishes strict from operational identity claims, and sketches testable implications in cognition, ethics, and AI-era identity systems.

Keywords: consciousness, semi-sameness, identity, operational identity, strict identity, persistence, equivalence relations, invariants, tolerance, binding, perception, memory, narrative, agency, responsibility, repair, forgiveness, categorization, auditability, artificial intelligence, centaur cognition, unrepeatable events.

A collaboration with GPT-5.2-Thinking. 42 pages.CC4.0.

Outline

1. Thesis and Framing
1.1 The non-repeating universe and the problem of sameness
1.2 The proposal: consciousness as an equivalence-maker
1.3 What “semi-sameness” means and does not mean
1.4 Claim scope and bracketing

2. The Ground Layer: Uniqueness Without Recurrence
2.1 Events versus descriptions: why repetition is descriptive, not ontological
2.2 Worldlines and continuity as the substrate of persistence talk
2.3 Why “same place” is always frame-relative
2.4 Repeatable experiments as repeatable procedures

3. Semi-Sameness as a Cognitive Construction
3.1 Feature extraction: invariants worth tracking
3.2 Tolerance assignment: eps depends on stakes
3.3 Binding: from features to objects and agents
3.4 Continuity chains: how “the same” persists over time
3.5 Where the construction fails: signature breakdowns

4. The Self as a Privileged Equivalence Class
4.1 Selfhood as owned continuity, not retained essence
4.2 Memory and narrative as re-identification engines
4.3 Responsibility across instants: accountability without rewind
4.4 Fission and duplication: why “same self” becomes underdetermined

5. Ethics Under Semi-Sameness
5.1 Irreversibility: moral weight in a one-pass world
5.2 Repair and forgiveness: restoration without erasure
5.3 Truth-labeling discipline: typed identity claims in moral speech
5.4 Dignity of the particular: resisting coercive categorization

6. Formal Discipline Without Mysticism
6.1 Minimal schema: id(t) as an index under a scheme S
6.2 Practical sameness: ~=_eps,F as equivalence under features and tolerance
6.3 Typed identity statements: strict versus operational
6.4 Why completeness is unavailable and how to be honest about it

7. Testable Implications and AI-Era Stakes
7.1 Cognitive predictions: binding failures and continuity disruptions
7.2 Institutional predictions: audit trails as continuity proofs
7.3 AI predictions: machine-made semi-sameness and error propagation
7.4 Design principles: auditability, humility, and repair orientation

8. Objections and Replies
8.1 “This is trivial: everything is unique anyway.”
8.2 “Equivalence is all we need; strict identity is idle.”
8.3 “Uniqueness destroys the self.”
8.4 “Structuralism is enough; haecceity is spooky.”

9. Conclusion: Uniqueness as Ground, Semi-Sameness as Tool
9.1 Restate the thesis and mechanism claim
9.2 What is gained: cleaner metaphysics, ethics, and language
9.3 What is risked: rhetorical inflation and overreach
9.4 Forward work: case studies and a disciplined Reference section

References

Art

1. Thesis and Framing

The motivating pressure behind this paper is simple to state and hard to metabolize: in a universe governed by spacetime, there are no repeated events. Nothing returns to the identical spacetime point with the identical total context. Yet human life is saturated with "again." We wake up in what we call the same bedroom, drink what we call the same coffee, talk to what we call the same person, and hold ourselves to what we call the same commitments. Our language, our institutions, and our moral intuitions operate as if sameness were a native property of reality.

This paper treats that "as if" not as an error to be scolded, but as a phenomenon to be explained. The central idea is that sameness is not discovered in the world as strict identity; it is constructed by cognitive systems as a workable, tolerance-based layer of equivalence. The paper names that layer semi-sameness and proposes that consciousness is the mechanism by which a finite agent produces and inhabits it. "Consciousness," on this approach, is not merely a private glow of experience; it is an active interface operation that compresses unrepeatable tokens into stable handles for action, memory, responsibility, and care.

The task of this opening section is to lay the conceptual rails. We will (i) state the problem in its strongest form, (ii) propose the mechanism claim precisely, (iii) define semi-sameness by contrast with nearby confusions, and (iv) bracket what is not being claimed so the thesis does not inflate into metaphysical overreach.

1.1 The non-repeating universe and the problem of sameness

Even before any philosophy, ordinary observation already hints that the world does not "come back" in a literal sense. The sunlight angle shifts, the body changes, the context drifts. But the stronger point is structural: events are indexed by spacetime and by their embedding in a causal history. To say an event repeats is not merely to say something similar happens; it is to claim the same event occurs again. Under the usual understanding of spacetime, that is not how events work. An event is a token with a location in history. Once it happens, it becomes part of the world's fixed past and cannot be re-entered.

And yet sameness is indispensable. Without it, there would be no stable reference and no reliable coordination. "That cup" would become an ungraspable blur of momentary appearances. "My friend" would dissolve into an untrackable series of social encounters. "I promised" would have no binding force because the promiser and the later self would not be linkable as the relevant agent.

So we face a tension that is not primarily scientific but practical: the world supplies novelty at the token level, while life requires stability at the handle level. The brain must reconcile these. It

must treat distinct instants as “the same enough” in order to act. If it does not, agency collapses into paralysis: no object is re-identifiable, no plan can be executed, no responsibility can persist.

The paper’s opening stance is that this reconciliation is not optional and not merely linguistic. It is a deep functional requirement. Any system that survives in time must implement a method for re-identification under change. That method is not strict identity. It is structured equivalence.

1.2 The proposal: consciousness as an equivalence-maker

The core proposal is a mechanism claim:

Consciousness is the integrated operation that constructs and maintains actionable equivalence classes over an unrepeatable event stream.

This is deliberately framed as an operational thesis rather than a metaphysical declaration about what consciousness “is” in its ultimate nature. The claim is about what consciousness does: it produces semi-sameness.

The mechanism has three components.

First, extraction of invariants. The system identifies features that remain stable enough across variation: shape, gait, voice timbre, location cues, social roles, causal regularities, norms. These invariants are not perfectly stable; they are stable relative to the agent’s needs and environment.

Second, tolerance assignment. The system sets an implicit epsilon, a “difference budget,” for what counts as the same in the present context. The epsilon tightens when stakes are high (surgery, weapons, contracts) and loosens when stakes are low (recognizing a friend from far away, finding your car in a parking lot). This is why sameness is not a single relation but a family of relations.

Third, binding and ownership. The system binds features into unitary objects and agents and assigns continuity: this object persists, this person persists, this self persists. The binding is not arbitrary. It is constrained by causal continuity and by predictive success: the bindings that work are reinforced.

In this view, consciousness is not just “what it is like” to see the cup. Consciousness is the act of making a cup into a stable object of experience rather than a flicker of sensations. It is the act of making “me” into a stable agent rather than a disconnected set of impulses. Subjectivity is not removed; it is reinterpreted as the lived interior of an equivalence-making process.

This proposal also yields a clear contrast with purely subpersonal processing. Many computations can occur without entering the integrated equivalence layer. The distinctive claim

here is that consciousness is where semi-sameness becomes available as a unified handle: not merely feature detection, but identity-like binding under tolerance, with implications for memory, intention, and moral ownership.

1.3 What "semi-sameness" means and does not mean

Semi-sameness is the name for the middle layer between strict uniqueness and naive sameness. It is "semi" because it is not identity, but it is not arbitrary either.

What it means:

Semi-sameness is a tolerance-scoped equivalence relation over token-states, grounded in invariants and continuity criteria, used for action and coordination.

It produces the ordinary judgments that make life possible: the same cup, the same person, the same promise, the same problem. These judgments are compressions. They ignore vast amounts of microdetail, and they must do so. The world is too detailed; the agent is finite.

What it does not mean:

It does not mean events repeat. Semi-sameness is not a denial of uniqueness at the event level. It is a practical response to it.

It does not mean "everything is just subjective." Semi-sameness is constrained by reality because failed equivalence classes lead to failed predictions and failed actions. The constraints are ecological: the environment punishes bad sameness.

It does not mean strict identity is achieved by better measurement. Even if measurement became perfect, strict identity across distinct instants would still not be obtained, because token-distinctness is not only about descriptive content. It is also about position in history and causal embedding.

It does not mean consciousness is magic. "Equivalence-maker" is not a supernatural claim; it is a functional description. The paper does not require any specific theory of mind to accept the mechanism claim. A reader can interpret "consciousness" through different models (global workspace, predictive processing, integrated information, higher-order thought) while still using the semi-sameness framework as a way to discipline identity talk.

Semi-sameness therefore aims to be a clarifying construct, not a metaphysical replacement religion. It names what we are already doing, then makes it explicit enough to evaluate.

1.4 Claim scope and bracketing

To keep the paper honest, we bracket several temptations.

First, we bracket ultimate ontology. The paper does not claim to solve what consciousness ultimately is or how qualia arise. It claims a functional role for consciousness in constructing re-identifiable handles.

Second, we bracket totalizing physics claims. The paper uses the non-recurrence intuition of spacetime as a framing constraint. It does not claim that physics proves the entire metaphysical picture. The argument is: given unrepeatability at the event level, explain how sameness works in practice.

Third, we bracket theological or mystical interpretations. Readers may choose to interpret particularity as spiritually significant, but the paper does not need that to run. It stays within conceptual analysis and mechanism talk.

Fourth, we bracket strong claims about AI consciousness. The paper will discuss machine-made semi-sameness in later sections as a functional construction (classification and re-identification at scale). Whether machines are conscious in the phenomenological sense is not required for the core argument. The key point is that systems can manufacture semi-sameness, and human consciousness is one such system. The question of whether AI has subjective interiority is acknowledged but not settled here.

Finally, we bracket the moral conclusion that “sameness is always oppressive” or “sameness is always good.” The paper’s stance is balanced: sameness is necessary, but it can be coerced, counterfeited, or weaponized. The ethical question is how to manufacture sameness with truth-labeling discipline and auditability, especially when stakes are high.

With these brackets in place, the thesis becomes tight: consciousness is proposed as the mechanism that constructs artificial semi-sameness, enabling agency in a universe of unrepeatable events. The remainder of the paper will elaborate the mechanism, formalize the typing discipline, and show where the framework changes how we reason in ethics, institutions, and AI-era systems.

2. The Ground Layer: Uniqueness Without Recurrence

The thesis of this paper does not require heavy astrophysics, but it does require a disciplined starting point: the world does not hand us literal repetition of events. What it hands us is a stream of unique tokens, and what we call repetition is, almost always, repetition of description: the reapplication of a template to a new token. This section lays down that ground layer so that

“semi-sameness” is not smuggled in as a metaphysical claim about the world, but recognized as a construction layered over unrepeatable occurrences.

The point is not to deny practical recurrence. Days recur. Seasons recur. Habits recur. Experimental outcomes recur. But what recurs is a pattern class, a procedure, a family resemblance across many distinct instants. The underlying token-events remain distinct, and that distinction matters precisely when people attempt to use “same again” language as if it licensed stronger claims than pattern similarity.

2.1 Events versus descriptions: why repetition is descriptive, not ontological

An event, in the strict sense used here, is a token: a particular occurrence at a particular instant embedded in a particular causal context. It is not merely “something like that happened.” It is that happening. Once it occurs, it belongs to the world’s history as a unique node.

A description, by contrast, is a type-level construct. It is a sentence, a model statement, a category label, a measurement report. Descriptions are repeatable because they are abstract. You can reuse the same words, the same model, the same label for many different tokens. That is their function: they compress complexity by ignoring detail.

This distinction is the simplest source of the illusion of recurrence. We are trained by language to mistake type reapplication for token return. We say “another sunrise” and forget that it is another token event. We say “the same argument again” and forget that it is a new event that resembles a previous one along salient dimensions. We say “this is the same as last time” and smuggle in a causal inevitability that the evidence may not support.

The paper’s discipline is to keep event and description separate.

An event is not the description of the event.

A repeated description does not entail a repeated event.

A repeated procedure generates a family of distinct events that can be described with similar language.

Once this is internalized, “semi-sameness” becomes intuitive. Semi-sameness is precisely the systematic reuse of descriptions and identity-handles across distinct events. It is what minds do when they construct equivalence classes: they treat multiple token-events as members of a type for purposes of prediction and action. The ground layer is therefore not “nothing repeats” in a rhetorical sense; it is “only descriptions repeat with perfect fidelity; events do not.”

2.2 Worldlines and continuity as the substrate of persistence talk

If events are unique, how does anything persist? The answer is continuity structure. In physical terms, persistence is not a thing reappearing; it is a connected history. The minimal formal image is a worldline or worldtube: a continuous trajectory through spacetime that links a sequence of events associated with a system.

The paper does not insist on any particular metaphysical interpretation of worldlines. It uses them as an explanatory scaffold: persistence talk becomes intelligible when we treat "the same object" as shorthand for "a chain of linked token-states connected by lawful continuity." This is the same move made in the first paper, but here it plays a different role. It is not the centerpiece; it is the ground layer upon which consciousness performs semi-sameness.

Continuity matters because it gives identity-talk something non-mystical to latch onto. When we say "the same cup," we are not claiming strict identity across instants. We are claiming that the present cup-token is connected to the prior cup-token by a continuity chain: causal continuity of matter, continuity of location relative to its environment, continuity of function and appearance, continuity of tracking by an agent.

This is also why the paper emphasizes that consciousness constructs *equivalence classes over continuity chains*. The chain is the substrate; the "same cup" handle is the cognitive compression. The world provides continuity; the mind manufactures the stable label.

Continuity talk also clarifies ambiguous cases. When continuity splits (duplication, fission), "the same" becomes underdetermined because the chain relation becomes one-to-many. That does not mean persistence is false; it means our operational identity criteria need to be declared. The ground layer gives the conceptual infrastructure for that declaration.

2.3 Why "same place" is always frame-relative

"Same place" feels like the simplest kind of sameness, and that is exactly why it is a useful diagnostic. People assume "place" is an intrinsic container in which events occur, and thus that returning to a place might be returning to a repeated event-context. But "place" is not absolute in the way ordinary language suggests. It is defined relative to a coordinate choice, a reference frame, and a scale.

Even in everyday life, "same place" is a compression. When you say you are standing in the same spot as yesterday, you mean "within a tolerance region relative to some stable landmarks." You do not mean the identical physical embedding: the Earth has rotated, the Earth has moved in its orbit, the solar system has moved, the gravitational environment has shifted, the air molecules are different, your body is different. The stability you rely on is relational: the

floorboards, the room boundaries, the landmarks that define the coordinate chart you are using.

This is the point: “place” is an operational concept. It is a constructed frame for coordinating action. It is a human-scale equivalence class over micro-different physical situations.

So “same place” becomes a clean example of semi-sameness before we even talk about consciousness explicitly. A conscious agent chooses a frame (usually implicitly), chooses a tolerance, and then treats a region of distinct token positions as one “place.” The usefulness of places is real, but the identity claim is operational.

This frame-relativity is not a philosophical trick. It is a reminder that much of what feels like solid sameness is already a tolerance-scoped construction. The world does not deliver places as absolute identity objects. It delivers relations that we stabilize into places.

2.4 Repeatable experiments as repeatable procedures

Scientific repeatability is sometimes mistaken as a refutation of uniqueness: if experiments can be repeated, aren’t events repeating? The framework answers: what repeats is the procedure and the pattern of results under controlled conditions, not the token event itself.

An experiment is a protocol: a designed set of operations intended to isolate variables and generate observable outcomes. When a lab repeats an experiment, it is generating a new family of token events engineered to be comparable. The comparability is the achievement. It requires boundary-setting, calibration, and a shared descriptive language. The repeated experiment is not the same event; it is a new event crafted to fall within a narrow equivalence class.

This is why repeatability is compatible with uniqueness and, in fact, depends on it. If events literally repeated, there would be nothing to test. Repeatability is meaningful precisely because each run is a distinct occurrence whose similarity is non-trivial. The discipline of science is the discipline of manufacturing semi-sameness under explicit tolerances: controlling error, specifying measurement resolution, and defining what counts as “the same result” within confidence bounds.

So scientific practice becomes an ally of this framework. Science is not built on strict identity. It is built on typed equivalence: same protocol, same measured outcome within tolerance, same statistical pattern across runs. Science is therefore a mature cultural technology of semi-sameness. It demonstrates, at scale, how powerful and honest operational sameness can be when its tolerances are declared and its failure modes are tracked.

This also sets up later sections on AI and institutions. Machine learning systems, identity systems, and bureaucratic classifications are also equivalence-makers. The question is whether their tolerances are disciplined, whether their linkage criteria are auditable, and whether their

operational sameness is allowed to masquerade as strict identity in contexts where people get harmed.

With this ground layer in place, the paper can now make its central move cleanly: consciousness is proposed as the organism-scale equivalent of the laboratory's repeatability discipline. It is the mechanism that manufactures "same enough" handles over unrepeatable events so that agency is possible.

3. Semi-Sameness as a Cognitive Construction

The paper's core claim is that semi-sameness is not found ready-made in the world; it is constructed by cognitive systems as a survival-grade interface. The world presents an unrepeatable stream of token-events. The agent must decide what to treat as stable, what to treat as change, and what to treat as "the same thing" across time. Semi-sameness is the solution to that decision problem: a scheme for mapping unique instants into operationally reusable handles.

This section describes the construction as a pipeline of functions rather than as a metaphysical mystery. The language will remain high-level, but each component is meant to be conceptually sharp: invariants, eps, binding, continuity, and breakdown. Together they explain why sameness is both necessary and dangerous: it is the tool that makes agency possible, and it is the tool that can misfire into illusion, coercion, or moral error when its assumptions go unexamined.

3.1 Feature extraction: invariants worth tracking

Semi-sameness begins with a question: what in the flux is worth tracking? A finite agent cannot track everything. So the first operation is feature extraction: selecting aspects of the sensory and internal stream that are stable enough, informative enough, and actionable enough to serve as re-identification anchors.

Invariants are not absolute sameness. They are stable patterns under transformation: features that persist under changes in viewpoint, lighting, distance, occlusion, and time. A face is a near-invariant across expressions; a voice has a near-invariant timbre across sentences; a familiar room has near-invariant geometry across small rearrangements.

But "invariant" here is broader than perception. There are also social invariants: role labels, status cues, norms of interaction, reputational patterns. There are moral invariants: promises, commitments, obligations, taboos. There are causal invariants: if I release the cup, it falls; if I anger this person in this way, that consequence often follows.

The key is that invariants are chosen under pressure. The environment rewards certain compressions. If you fail to extract the right invariants, you die, you lose social standing, you miscoordinate, or you suffer preventable harm. So the invariants are not arbitrary. They are shaped by ecological constraint and by reinforcement.

This also explains why different agents can disagree about sameness. They may be tracking different invariants. A mechanic identifies a car by engine sound patterns and maintenance history; a child identifies it by color and shape; an insurer identifies it by VIN and paperwork. These are not competing realities; they are different feature sets chosen for different tasks.

So semi-sameness begins as a selective projection: the world is mapped into a reduced feature space. That projection is the first act of "making the world usable."

3.2 Tolerance assignment: eps depends on stakes

Here, *eps* (short for *epsilon*) denotes the tolerated degree of difference under which two distinct token-states are treated as operationally the same for a given purpose. After selecting features, the agent must decide how similar two token-states must be to count as "the same" under those features. This is where eps enters. Eps is not a single number; it is a family of tolerances that govern equivalence decisions.

In low-stakes contexts, eps can be loose. You recognize your friend from a distance even if the face features are degraded. You find your car in the lot even if you cannot see the license plate. You call it the same coffee even if today's brew is slightly different. Loose eps supports fluidity and speed.

In high-stakes contexts, eps tightens. A pharmacist cannot accept "close enough" on a dosage label. A pilot cannot accept "close enough" on altitude. A court cannot accept "close enough" on identity when liberty is at stake. Tight eps supports safety and justice.

This stakes-dependence is the hidden heart of semi-sameness. It shows that sameness is not a binary property; it is a regulated decision. And it reveals why arguments about "same" often feel like moral fights: the parties are using different eps values without naming them. One side is treating the situation as low stakes; the other is treating it as high stakes.

Eps also varies with context and stress. Under fear, the mind may tighten eps in some dimensions ("is that threat the same as last time?") and loosen it in others ("everything dangerous is the same"). Under fatigue, eps may loosen unintentionally, producing errors of conflation. Under ideology, eps can become distorted: categories become too broad or too narrow to serve truth.

So the paper's mechanism claim has an ethical edge: consciousness does not merely assign eps; it assigns eps in ways that can be trained, manipulated, or corrupted. Semi-sameness is

therefore not only cognitive; it is normative. The tolerance setting is part of the moral life, because it determines who is treated as "the same kind of person," what counts as "the same offense," and what counts as "the same repair."

3.3 Binding: from features to objects and agents

Feature extraction and tolerance assignment still do not yield "things." They yield similarity judgments in a feature space. Binding is the further operation that turns those judgments into coherent objects and agents: units that can be tracked, owned, and acted upon.

Binding answers questions like:

Which features belong together as one object rather than many?

Which sensory streams are caused by the same source?

Which actions belong to one agent rather than to noise or external force?

In perception, binding is what produces objecthood: the cup is a unit, not a collection of color patches and edges. In action, binding is what produces agency: the movement was my doing, not merely something that happened. In social life, binding is what produces personhood: the voice, the face, the name, and the history bind into "this person."

Binding is also where consciousness becomes most salient in lived experience. The conscious scene presents already-bound objects, already-bound agents, already-bound meanings. The claim here is that this is not merely a passive display. It is the output of an integrative operation that makes the world appear as a world of stable units rather than as raw flux.

Binding is not infallible. It can misbind. It can bind separate sources into one (hallucinated unity), or split one source into many (fragmentation). It can bind the wrong history to the wrong face (misrecognition). It can bind guilt to the wrong agent (scapegoating). These failures are not peripheral; they are diagnostic of how the system works. They show that "identity" is an active construction, not a metaphysical given.

Binding also connects directly to the earlier claim that consciousness is an equivalence-maker. The equivalence class only becomes practically usable when it is bound into an object that can be referred to. Semi-sameness without binding is like similarity without naming: it cannot ground action or obligation.

3.4 Continuity chains: how "the same" persists over time

Once objects and agents are bound, the system must maintain them across time. Continuity chains are the cognitive analog of worldlines: a sequence of token-states linked under tracking assumptions.

A continuity chain is a running hypothesis:

The object I saw at t1 is linked to the object I see at t2 by a plausible continuity path under my model of the world.

The chain is updated by prediction and correction. If the object disappears behind a wall and reappears, the system infers continuity. If the reappearance violates constraints (wrong size, impossible location), the chain is broken and a new chain is formed.

This is how semi-sameness becomes persistence talk. “Same cup” means: I have maintained a continuity chain that links these token-states under my tracking criteria. The claim is operational. It is also corrigible: if evidence breaks the chain, the sameness claim should be revised.

Continuity chains also operate for the self. The “I” is a privileged chain maintained across memory, bodily continuity, and agency attribution. The self-chain is reinforced by ownership: actions are tagged as mine; consequences are inherited; plans are carried forward. The sense of personal identity is, on this view, the lived interior of a continuity chain that has become stable and socially reinforced.

This chain model clarifies why certain disruptions are so existentially unsettling. If the continuity chain for objects breaks, the world feels unreal. If the continuity chain for the self breaks, the self feels unreal. The paper will not medicalize this; it uses the point conceptually: persistence is not a metaphysical substrate, it is a maintained chain under constraints.

3.5 Where the construction fails: signature breakdowns

If semi-sameness is a construction, it must have characteristic failure modes. These failures are not merely errors; they are signatures that reveal the moving parts.

Several families of breakdowns follow directly from the pipeline described above.

Invariant failure. The agent extracts the wrong features or fails to stabilize the right ones. This yields fragile re-identification. Everything feels new and untrackable, or everything collapses into crude stereotypes. Invariant failure can come from sensory degradation, stress, or learned bias.

Eps miscalibration. The tolerance becomes too loose or too tight for the stakes. Too loose yields conflation: different people treated as the same, different situations treated as equivalent threats, different actions treated as morally interchangeable. Too tight yields fragmentation: the same person feels like a stranger after minor change, the same institution feels illegitimate after a small shift, the self feels discontinuous across normal development.

Binding errors. Features are grouped incorrectly. This yields classic misattributions: the wrong voice to the wrong face, the wrong intention to the wrong act, the wrong causal source to a percept. In social terms, binding errors are the engine of scapegoating and of certain kinds of paranoia: the system binds multiple events to one hidden agent because the narrative fit is too strong.

Continuity chain breaks and false continuities. The system either breaks a chain when it should not (producing disorientation) or maintains a chain when it should not (producing deception). False continuity is especially important ethically and institutionally: a forged identity document is a forced continuity chain; propaganda is a forced continuity narrative; counterfeit repair is a forced continuity claim without underlying change.

Narrative override. A special failure mode occurs when narrative templates dominate sensory and evidential updating. The system stops building chains from evidence and starts forcing evidence into a pre-made chain. This is how "it's the same as last time" becomes a trap: it collapses uniqueness into a template and disables revision. The paper treats this as one of the main hazards of semi-sameness: it is necessary, but it can become totalizing and coercive.

These breakdowns are not just cognitive curiosities. They are the bridge to later sections on ethics and AI. When institutions and machine systems manufacture sameness, they can commit the same failures at scale: miscalibrated tolerances, binding errors, false continuities, and narrative overrides disguised as objective classification. The framework therefore offers not only a theory of consciousness but also a diagnostic language for modern equivalence machinery.

With the cognitive construction pipeline now in view, the paper can proceed to its most charged application: the self. If consciousness constructs semi-sameness, then the self is not an exception; it is the most intense and morally consequential instance of the same construction.

4. The Self as a Privileged Equivalence Class

If semi-sameness is the mind's method for making a usable world, then the self is the mind's most consequential product. The self is not merely another object among objects. It is the center of control, the bearer of responsibility, the locus of suffering and joy, the point where continuity becomes morally charged. That is why identity debates become heated when they touch the self: people are not merely arguing about metaphysical labels; they are defending the conditions under which promises bind, guilt makes sense, love persists, and repair is possible.

This section proposes a disciplined re-description: the self is a privileged equivalence class maintained by owned continuity. The self is not a retained essence that stays numerically identical across time. It is a continuity chain that the system treats as "mine," and that others

treat as the bearer of accountability. That chain is stabilized by memory, narrative, bodily persistence, social recognition, and the inheritance of consequences. The framework therefore aims to preserve everything that matters about the self (agency, dignity, accountability) while refusing metaphysical overreach (event recurrence, strict identity across time).

4.1 Selfhood as owned continuity, not retained essence

Ordinary talk often treats the self as a thing that persists: a substance-like core that remains the same while experiences and circumstances change. That picture is powerful because it matches a felt intuition of "I am here again." But under instantaneous uniqueness, the "again" cannot mean strict identity at the level of events or token-states. The more disciplined reading is that "I" refers to a maintained continuity chain that is marked by ownership.

Ownership is the crucial distinction. Many continuities exist in the world: the continuity of a river channel, a business, a hurricane, a tool. The self is the continuity that the system treats as the subject of agency and the heir of consequences. It is the chain for which the system says, implicitly, "this is the one that acts and is acted upon as me."

Several features make this continuity privileged.

Agency coupling. The self-chain is tightly coupled to decision and action. Plans formed at one time are carried forward as commitments. Errors become lessons. The system tracks not just what happened, but what it did.

Consequence inheritance. The self-chain is where effects accumulate. A broken bone continues to matter. A reputation continues to matter. A promise continues to matter. This inheritance is not metaphysical; it is causal and social. It is nevertheless what gives the self its temporal thickness.

Normative addressability. The self-chain is the one that can be addressed by praise, blame, request, forgiveness, and obligation. "You" is a social pointer to a continuity chain.

Bodily anchoring. For humans, the continuity chain is anchored in a body that is itself a stable physical process. The body is not a retained essence, but it provides an unusually strong continuity substrate, especially compared to many artifacts.

So selfhood, in this view, is not a hidden pearl. It is a stabilized chain plus a social and internal practice of ownership. The metaphysical shift is subtle but profound: it moves the self from "a thing that stays the same" to "a history that stays linked."

This also changes how we understand growth. If the self is not a retained essence, then change is not a threat to identity; it is the normal mode of identity. The self persists by maintaining enough linkage to remain accountable and intelligible, not by remaining unchanged.

4.2 Memory and narrative as re-identification engines

If the self is a continuity chain, then memory and narrative are the machinery that maintains it. They do not merely store the past; they perform re-identification. They take the unrepeatable and make it present enough to bind the chain.

Memory provides internal linkage. It supplies tokens of past experience as present representations. By carrying these representations, the system can say: the one who remembers is the heir of the remembered act. This is why memory loss can feel like identity loss: it weakens the internal continuity proof.

Narrative provides social and interpretive linkage. It compresses a life into a story that can be told to others and told to oneself. Narrative selects which events count as defining, which changes count as growth, and which ruptures count as breaks. Narrative is therefore not merely decorative. It is a chain-maintenance tool.

This dual engine explains several familiar phenomena.

First, identity feels stronger when memory and narrative cohere. A coherent life story makes the self feel continuous. This is not necessarily truth; it is a form of internal auditability. When the story aligns with records and relationships, continuity feels earned.

Second, identity feels threatened when narrative fractures. Contradictory stories about oneself can weaken ownership: the system does not know which chain it is maintaining. This is why confessional practices, therapy-like reflection, and even ordinary journaling can feel stabilizing: they are attempts to re-align narrative with lived continuity.

Third, identity is socially negotiated. Other people participate in maintaining your continuity chain by recognizing you, recalling shared history, and holding you to your commitments. This can be beautiful (love as continuity stewardship) or coercive (refusing to allow someone to change by insisting they remain “the same person” in a punitive sense).

So memory and narrative are not just contents in consciousness; they are active operators that manufacture semi-sameness of the self across time.

This also highlights a moral hazard: because narrative is compressive, it can counterfeit continuity. A person can craft a narrative that implies they have repaired when they have not, or that implies they are the same as their former self in ways that excuse harm. The framework therefore treats narrative as a tool that must be audited against action and consequence, not as a sovereign authority over identity.

4.3 Responsibility across instants: accountability without rewind

The strongest stress-test of any identity theory is moral responsibility. If the self is only a chain of unique instants, why should later instants be accountable for earlier ones? The framework answers: because responsibility binds to continuity of agency and consequence inheritance, not to metaphysical strict identity.

Responsibility is a typed claim. It asserts that the present agent-state is the rightful heir of the earlier act because the chain linking them satisfies relevant criteria.

Causal continuity. The present state exists in the causal aftermath of the earlier act. It benefits from it, suffers from it, or stands in its wake.

Agency continuity. The earlier act issued from the same decision-making system: the same capacities, values, and controls that the present state inherits, unless a genuine rupture is shown.

Normative continuity. The agent participated in a norm-governed context where promises and obligations were made intelligible and binding.

This preserves accountability while avoiding metaphysical fiction. The system does not need to claim it can redo the event or return to the same moment. It needs to claim that the chain is continuous enough that obligations carry forward.

Irreversibility therefore increases moral seriousness rather than decreasing it. If the event cannot be undone, then responsibility becomes forward-directed: the responsible agent is the one who must repair, restrain future harm-patterns, and bear consequences. Accountability is not about metaphysical sameness; it is about stewardship of the chain.

This also clarifies why certain disruptions mitigate responsibility. If the chain's agency continuity is damaged (coercion, severe impairment, radical discontinuity), the responsibility claim must be adjusted. The framework does not invent these mitigations; it explains their rationale.

So "accountability without rewind" becomes a central ethical payoff: you can maintain moral continuity in a world of unique instants by grounding responsibility in demonstrable continuity criteria rather than in an unargued essence.

4.4 Fission and duplication: why "same self" becomes underdetermined

The hardest cases arise when continuity chains branch or are duplicated. These cases are often presented as science-fiction puzzles, but their philosophical function is to expose hidden assumptions in ordinary identity talk.

Fission is a branching continuity: one earlier self-state leads to two later self-states, each with strong continuity in memory, character, and agency structure. Duplication is the creation of another self-like state with similar internal organization and memories, whether by copying information or replicating structure.

Under a retained-essence model, branching is paradoxical because strict identity is assumed to be one-to-one. Under the chain model, branching is not paradoxical; it is underdetermination. The question “which one is me?” has no single answer unless we stipulate extra criteria.

Several responses are possible within the framework, and the key is that each response corresponds to a different operational identity rule.

Bodily continuity rule. “Me” follows the continuity of a particular biological organism. Branching in psychological continuity does not override bodily anchoring.

Psychological continuity rule. “Me” follows memory and character continuity. Under branching, “me” becomes a relation that can have multiple successors, which breaks one-to-one identity but preserves what matters psychologically.

Normative-ownership rule. “Me” follows the chain that inherits social and legal accountability structures. Under branching, institutions may choose one successor as the accountable continuation for practical reasons, or split accountability according to causal roles.

The paper’s position is not to legislate a single answer. Its position is to force honesty: “same self” is not a primitive fact in branching scenarios. It becomes a typed claim whose truth depends on which continuity criteria we choose to privilege.

This is not a weakness. It is a revelation about ordinary life. Even without science-fiction branching, we already encounter softer versions of underdetermination: radical personality change, memory loss, addiction recovery, religious conversion, trauma. People ask, “am I the same person?” and the answer depends on which continuity dimensions are morally relevant: bodily persistence, agency control, relational commitments, character invariants.

So fission and duplication do not refute the self. They show that selfhood is a privileged equivalence class maintained by continuity criteria, and that when continuity criteria diverge, identity judgments become context-dependent rather than metaphysically absolute.

This completes the central move of the paper. Semi-sameness is not merely a cognitive convenience applied to cups and chairs. It is the architecture by which a self becomes a self: a continuity chain owned as mine, stabilized by memory and narrative, capable of accountability without rewind, and honest about its limits when branching forces our hidden assumptions into the open.

5. Ethics Under Semi-Sameness

If semi-sameness is how consciousness makes a usable world, then ethics is where that usability becomes accountable. Ethics is not optional commentary layered on top of cognition. Ethics is the governance of equivalence-making under stakes: which differences we are permitted to ignore, which continuities we are obligated to honor, and which sameness claims become harmful when treated as stronger than they are.

The central ethical problem is now visible. Semi-sameness is a necessary compression. Without it, agency collapses. But compressions can lie, not only by falsifying facts, but by silently changing their scope. They can collapse distinct people into one category, collapse distinct harms into "the same kind of thing," collapse distinct responsibilities into "everyone's fault," or collapse genuine repair into symbolic repetition. The ethical task is therefore not to abolish sameness language but to discipline it: to insist that "same" is a typed claim with declared criteria and tolerances, especially when dignity, justice, and harm are at stake.

This section develops four consequences: irreversibility as moral weight, repair and forgiveness as forward construction, a truth-labeling discipline for identity claims, and the dignity of the particular as an ethical counterweight to coercive categorization.

5.1 Irreversibility: moral weight in a one-pass world

Instantaneous uniqueness implies irreversibility in the strong, everyday sense: you do not get the same moment back. The moral meaning of irreversibility is not merely that outcomes persist. It is that actions are unrepeatable as tokens. A cruel sentence is spoken once, in that context, with that vulnerability. A betrayal happens once, and the relational world after it is different forever. A neglected child does not get the same day re-issued with improved parenting.

Most moral evasions rely, implicitly, on denying irreversibility. People behave as if they can "try again," as if apology can null the act, as if a ritual can reset the ledger, as if time can be negotiated by declaration. Semi-sameness makes these evasions intelligible: the mind's equivalence machinery wants to treat "another day" as "the same day form," and then smuggle that form into moral accounting as if the harm can be overwritten.

The discipline here is severe but liberating: you cannot undo an event; you can only make the next one. That converts ethics from a fantasy of rewind into a craft of forward-making. It also elevates the moral weight of attention. If each encounter is a one-pass event, then the ethical life is not primarily a set of abstract principles; it is the fidelity with which you treat the particular moment and person before you.

Irreversibility also clarifies the difference between regret and responsibility. Regret is an internal attitude toward an unrepeatable act. Responsibility is the outward obligation to inherit consequences and to repair. In a one-pass world, regret without repair is a counterfeit moral motion: it rehearses an inner scene as if replay could substitute for restitution. Repair is the only morally meaningful "second chance," and it is not a return; it is a new event that addresses an old one.

Finally, irreversibility exposes a crucial asymmetry: the powerful often demand equivalence for their own benefit ("that was the old me," "that was a different time," "everyone does it"), while insisting on strictness for others. In a disciplined framework, you cannot select tolerances opportunistically. If you loosen eps to excuse yourself, you must be willing to loosen it in comparable cases, or admit you are not describing but bargaining.

5.2 Repair and forgiveness: restoration without erasure

Once irreversibility is accepted, repair and forgiveness become conceptually clean.

Repair is not erasure. It is the construction of new continuity that makes the damaged world more habitable. It is forward-directed work that takes the unrepeatable past as fixed and then acts to reduce its harm, restore capacity, and prevent recurrence of the harm-pattern.

Several elements typically belong to repair.

Acknowledgment. A truthful description of the unrepeatable act and its consequences, resisting the temptation to compress it into a generic category ("mistakes were made").

Restitution. Concrete acts that restore what can be restored, materially and relationally, without pretending full restoration is always possible.

Constraint change. Modification of the conditions that made the harm likely: habits, systems, incentives, access, supervision, social arrangements. This is the practical meaning of repentance: not mood, but changed constraints.

Memory stewardship. Preservation of the record in a way that honors the harmed and prevents narrative laundering. This does not mean obsession; it means refusing to let the past be rewritten by convenience.

Forgiveness is different. Forgiveness is not the cancellation of history. It is a decision about how the injured party will carry the history forward. It can be offered without denying truth and without waiving justice. In this framework, forgiveness is best understood as a controlled relaxation of certain equivalence operations in the future: the harmed person chooses not to treat every subsequent token-event as equivalent to the harm-event, even when the

resemblance cues are present. Forgiveness is therefore a discipline of not overgeneralizing the past into a totalizing template.

This matters because many social systems confuse forgiveness with erasure. They demand that the harmed compress the harm-event into a type ("it's over," "move on") and treat future events as if the harm never occurred. That is not forgiveness; it is coercive semi-sameness: forced equivalence between "forgiven" and "forgotten." The framework refuses that conflation. Forgiveness can be compatible with boundaries, memory, and even ongoing consequences. It is a gift of future openness, not a claim that the past is unreal.

So restoration without erasure becomes the ethical signature of semi-sameness done well: a person can choose not to let one event become the equivalence class for all future events, while still refusing any lie about what happened.

5.3 Truth-labeling discipline: typed identity claims in moral speech

Because ethics often turns on identity claims, the paper proposes a truth-labeling discipline for moral speech. The goal is not to change how people talk in casual life; it is to prevent high-stakes moral and institutional harm that arises when untyped "same" language is treated as metaphysical truth.

The discipline has four rules.

First, prohibit strict identity claims about events. Do not speak as if you can re-enter the moment, undo it, or "make it as if it never happened." Say instead what you mean: "I want to repair what I did," "I want to prevent this pattern," "I want to restore trust."

Second, type sameness claims about persons. When someone says "he's the same person," or "I'm not the same person," require the operational criteria. Do they mean continuity of body, continuity of agency, continuity of memory, continuity of character, continuity of accountability structures? Different claims answer to different evidence. Without criteria, "same person" becomes a rhetorical weapon: used to deny growth or to evade responsibility.

Third, declare tolerances when stakes are high. Moral disputes frequently hide eps conflicts. One party is using a loose eps ("close enough, let it go"), the other a strict eps ("this difference matters"). The ethical requirement is to make that explicit. What differences are permitted to be ignored, and why? Who bears the risk if the tolerance is wrong?

Fourth, separate descriptive equivalence from normative equivalence. Two acts can be descriptively similar and normatively different because of intention, capacity, and context. Conversely, two acts can be descriptively different and normatively equivalent in harm. The discipline demands that speakers specify which equivalence they are claiming: similarity of appearance, similarity of causal structure, similarity of moral weight, similarity of legal category.

These rules reduce manipulative compression. They protect against three common abuses.

Counterfeit apology. Words that mimic repair while avoiding acknowledgment, restitution, and constraint change.

Counterfeit renewal. Ritualized “fresh starts” that imply erased obligations.

Counterfeit blame. Category-based guilt that binds harm to the wrong agent by loose equivalence and lazy binding.

Typed moral speech is therefore not academic fussiness. It is a protection of dignity and justice in a world where equivalence-making is powerful and easy.

5.4 Dignity of the particular: resisting coercive categorization

Semi-sameness is necessary, but it carries a built-in danger: categorization can become coercive. When institutions or ideologies manufacture equivalence classes, they can flatten persons into types, and then treat the type as if it were the person. This is the ethical catastrophe of sameness: the replacement of the particular with the label.

The framework responds by making particularity morally salient. If events are unrepeatable and persons are continuity chains rather than interchangeable instances, then each person has a dignity that is not reducible to category membership. Categories remain useful, but they must be held as operational tools, not as total descriptions.

Resisting coercive categorization involves several disciplines.

Preserve token-level attention. Ethical perception begins with seeing the particular: this person, this harm, this context. Categories may guide, but they must not replace seeing.

Refuse identity laundering. Do not let narrative templates reclassify a person’s lived history into a convenient type that erases their truth. This includes both flattering and condemning templates.

Protect the right to change. Coercive sameness often takes the form of refusing to allow a person to become different: “you’re always like this,” “people like you never change.” The chain model allows change while maintaining accountability. It says: growth is real, but obligations persist; repentance is real, but it must be evidenced by constraint change.

Guard against machine-made coercion. Modern classification systems amplify coercive categorization by producing equivalence classes at scale and treating them as authoritative. The ethical demand is auditability: categories must be contestable, tolerances must be justified, and recourse must exist when the equivalence machine misbinds a life.

Under this view, dignity is not a sentimental add-on. It is the ethical counterweight to the cognitive necessity of compression. Because consciousness constructs semi-sameness to function, ethics insists that the construction remain humble: always aware that labels are tools, that persons are histories, and that the unrepeatable should not be crushed into an interchangeable type.

This completes the ethical arc of the paper's central claim. Consciousness manufactures semi-sameness so that life can proceed in a non-repeating world. Ethics governs that manufacture so that the world does not proceed by lies, erasures, and coercive equivalences. The next sections will formalize the discipline more explicitly and then extend it to institutional and AI-era equivalence systems, where semi-sameness is now produced not only by minds but by machines.

6. Formal Discipline Without Mysticism

"Semi-sameness" will be dismissed as a mood unless it can be spoken with discipline. Discipline does not require turning the paper into mathematics, but it does require that identity talk be made explicit enough to audit: What is being compared? Under what features? With what tolerance? By what linkage rule? This section provides a minimal formal scaffold that accomplishes that goal while avoiding metaphysical inflation.

The guiding principle is modesty. We do not posit hidden substances or "thisness particles." We use indices, relations, and typed claims. These are not metaphysical replacements for lived reality; they are tools that prevent category mistakes. They let us say precisely what kind of sameness we mean and what kind we cannot mean in a universe of unrepeatable events.

6.1 Minimal schema: id(t) as an index under a scheme S

Begin with the simplest act: reference. We want to refer to "this token-state at this instant" without implying that the reference is a metaphysical essence.

Let x denote an entity-track candidate: a person, object, system, institution, or process we intend to talk about as persisting. Let t denote a time parameter along some chosen ordering or clock relevant to the discussion. At each t, we have a token-state x(t), meaning "x as it is at t," not as a timeless essence.

We then define an identifier:

id_x(t; S)

Read: the identifier of x at time t under scheme S.

Scheme S is crucial. It represents the tracking and description conventions we are using. Different schemes emphasize different features, boundaries, and linkage rules. A scheme might specify:

What counts as the system boundary (what is included as x, what is environment).

Which features are tracked (geometry, mass, serial number, memory state, legal charter).

What counts as continuity (causal propagation, record linkage, bodily persistence).

What resolution is assumed (measurement precision, record granularity).

So id_x(t; S) is not "the soul of x." It is a pointer created by a tracking practice. It can be implemented as a coordinate address in a model, a database key, a legal identifier, or an internal cognitive handle. The key point is that it is *scheme-relative*.

This already blocks a common metaphysical mistake: treating an identifier as if it were an intrinsic property of the thing rather than a tool for referencing it under a practice.

Within the paper's thesis, instantaneous uniqueness is captured by the idea that token-states are distinct nodes in history. Even if x persists in the operational sense, x(t1) and x(t2) are different token-states when t1 != t2. The id index therefore changes with t, even if the tracking scheme judges the two token-states as belonging to "the same" continuity chain.

This yields a clean separation:

Uniqueness applies to token-states.

Sameness applies to relations among token-states under a scheme.

6.2 Practical sameness: ~=_eps,F as equivalence under features and tolerance

We now formalize "same enough." Practical sameness is an equivalence relation (or at minimum a similarity relation) defined relative to a feature set and a tolerance.

Write:

a ~=_(eps, F) b

Read: a is practically the same as b within tolerance eps when comparing features F.

Here:

F is the set of features being compared. For a face, F might include geometry and voice cues. For a legal identity, F might include name, date of birth, and document signatures. For a scientific measurement, F might include instrument outputs and calibration conditions.

eps is the tolerance or slack allowed. It can be numerical (measurement error bounds) or procedural (accepted verification thresholds, standards of proof, acceptable deviation for function).

This expression is the formal heart of semi-sameness. It says: sameness is not primitive; it is defined by what you decide to measure and what differences you allow yourself to ignore.

Several consequences follow immediately.

Sameness is plural. Different choices of F and eps produce different equivalence classes. That is not incoherence; it is reality. “Same person” in friendship and “same person” in banking do not use the same F or eps.

Sameness is stakes-sensitive. High stakes demand smaller eps and stricter F. Low stakes permit looser eps and broader F.

Sameness is corrigible. If new evidence changes the relevant features or reveals that eps was too loose, the equivalence judgment must be revised.

This also clarifies the paper’s claim that consciousness is an equivalence-maker. The cognitive system is constantly choosing (often implicitly) both F and eps: which features to attend to and how much deviation to allow before it breaks the continuity chain and declares “different.”

6.3 Typed identity statements: strict versus operational

The paper’s biggest practical payoff comes from typing identity claims so that speakers cannot slide between meanings.

We define two “identity-like” relations:

=*{strict} and* ={op}

x =*{strict} y means: x and y are numerically identical tokens, the same token-state or the same event. Under instantaneous uniqueness, this relation is extraordinarily rare across time. For token-states, x(t1)* ={strict} x(t2) is false whenever t1 != t2.

x =*{op} y means: x and y are operationally identical under scheme S for a purpose P. This typically means they are linked by continuity rules and/or satisfy practical equivalence under* ~=(eps, F) in ways sufficient for the purpose.

So a more explicit form is:

x =_{op,S,P} y

This notation forces the hidden parameters into view: scheme and purpose.

Most everyday uses of “same” are operational, not strict. That is not a problem. The problem is untyped speech that pretends operational identity is strict identity when convenient, and then retreats to “it was only operational” when challenged. Typed claims prevent that drift.

This typing also clarifies moral speech. “I am not the same person” might mean:

My current state is not ~=_(eps, F) to my past state under features relevant to moral risk.

Or:

My agency continuity has changed because constraints changed.

It rarely means:

I have no continuity obligations.

Similarly, “you are the same person” might mean:

Your continuity chain is intact for accountability purposes.

It should not mean:

You are metaphysically frozen and cannot change.

So typing identity claims is not an academic flourish. It is a practical discipline that protects against moral evasion and coercive labeling.

6.4 Why completeness is unavailable and how to be honest about it

A common response to scheme-relativity is to demand a “complete description” that would settle identity absolutely: if we knew everything about a system at an instant, we could decide whether it is the same. The paper rejects this as both impossible in practice and confused in concept.

Completeness is impossible in practice because any real description has finite resolution, finite measurement, finite storage, and finite access. Every identification system—from perception to law—operates under limited features and bounded tolerances.

Completeness is confused in concept because it smuggles in the assumption that identity is exhausted by descriptive content. But even perfect content-matching would not automatically yield strict identity across time, because token-distinctness includes historical embedding: where in the causal structure the token occurs. Two states could match in every feature you record and still be different tokens with different provenance. Identity is not only “what it is like” in feature space; it is also “where it sits” in history.

So the book-keeping must remain partial. The ethical and epistemic requirement is not completeness. It is honesty about incompleteness.

Honesty has three components.

Declare features. Say what you are comparing. If you claim sameness without stating F, your claim is vulnerable to hidden bias.

Declare tolerance. Say how much difference you allow. If you claim sameness without eps, your claim is either rhetorical or naive.

Declare purpose. Say what the sameness is for. “Same person” for friendship is not “same person” for punishment. Without purpose, sameness becomes a weapon.

This discipline aligns with the paper’s broader posture: semi-sameness is unavoidable, but it must be governed. Formal minimalism provides that governance without mysticism. It supplies the vocabulary that keeps the thesis from inflating into a cosmic slogan and keeps moral and institutional speech from laundering operational equivalences into metaphysical certainties.

With this scaffold in place, the remaining sections can make sharper claims about testable implications and AI-era stakes. We can now ask, precisely, where consciousness should set eps, how institutions encode schemes S, how machine systems impose feature sets F, and where failures of binding and continuity become predictable—and preventable—forms of harm.

7. Testable Implications and AI-Era Stakes

A framework earns its keep when it changes what we expect to observe, how we diagnose failures, and what we build to prevent harm. The semi-sameness thesis does all three. It predicts characteristic breakdowns when the equivalence-making machinery is stressed or damaged. It predicts which institutional designs will better preserve continuity without coercion. And it predicts a new hazard in the AI era: machine-made semi-sameness that is powerful, opaque, and scalable, capable of propagating misbinding and coercive categorization far beyond the reach of any single human consciousness.

“Testable” here is meant in an ordinary, pragmatic sense. These are not predictions that require perfect instrumentation or final theories of consciousness. They are predictions that can be checked by patterns: by which errors cluster where, by how systems behave when constraints change, by whether certain design choices systematically reduce specific harms. In this sense, semi-sameness functions like a mechanism model: it identifies parts (features, eps, binding, continuity, ownership) and predicts failure signatures when those parts miscalibrate.

7.1 Cognitive predictions: binding failures and continuity disruptions

If consciousness is an equivalence-maker, then disruptions to consciousness or to its integration functions should produce signature errors in "sameness construction," even when lower-level sensory processing remains intact. The predicted failures are not random; they map onto the pipeline described earlier.

Prediction family A: object binding failures. When binding weakens, the perceptual world becomes less "object-like." The agent may still register features—colors, edges, sounds—yet have difficulty stabilizing them into enduring units that can be tracked. The subjective correlate, when present, is often described as fragmentation: the world feels disjointed or less real. The testable content is behavioral: the system's re-identification performance drops under occlusion, motion, or distraction, and it makes more errors of "new object" versus "same object" across time.

Prediction family B: false bindings under narrative pressure. When top-down templates dominate updating, the system binds multiple distinct events into a single presumed cause or agent. The testable marker is over-attribution: different occurrences are treated as "the same thing happening again" despite weak evidence. This can manifest as stereotyped threat generalization ("every similar cue is the same danger") or stereotyped social categorization ("every member of a group is the same"). The mechanism claim is that eps becomes too loose and feature sets become too coarse under stress or ideology, yielding coercive equivalence.

Prediction family C: continuity-chain breaks in self-representation. If the self is a privileged continuity chain, then disruptions to memory integration, ownership tagging, or agency inference should produce characteristic distortions in personal continuity judgments: difficulty in linking earlier commitments to present obligation, or difficulty recognizing one's own past actions as "mine." The point is not to medicalize; it is to specify what to look for: a decoupling between memory content and ownership, and between action outcomes and inherited responsibility.

Prediction family D: eps miscalibration with stakes. The framework predicts that under fatigue, fear, or novelty, the system's tolerance setting becomes less aligned with objective stakes. Errors cluster where the cost of misidentification is high: confusing similar-looking items, overgeneralizing from one negative event, misreading intentions. Conversely, in some cases fear may tighten eps too aggressively, producing over-fragmentation: the friend "doesn't feel like the same person," the home "doesn't feel like the same place," the self "doesn't feel continuous." These are testable as context-dependent shifts in re-identification thresholds.

All of these predictions can be evaluated without settling what consciousness "is" in an ultimate sense. They are predictions about the functional role attributed to consciousness: integrated equivalence-making with binding, tolerance, and continuity maintenance.

7.2 Institutional predictions: audit trails as continuity proofs

Institutions are externalized equivalence-makers. They manufacture operational sameness through identity documents, ledgers, case files, certifications, contracts, and procedural definitions. Under the semi-sameness framework, institutional legitimacy depends on whether these manufactured sameness claims are auditable and contestable rather than coercive and opaque.

The central institutional prediction is this:

Where stakes are high, systems that encode continuity as auditable chains will outperform systems that rely on static labels.

This prediction is checkable across domains.

In finance, chain-of-custody and transaction logs reduce disputes about ownership and responsibility.

In healthcare, auditable records reduce medication errors and enable accountability when errors occur.

In law, documented procedures and evidence chains reduce arbitrary identity assignments and reduce "because we said so" coercion.

In organizations, decision logs prevent post hoc narrative rewriting ("that's not what we meant") and therefore reduce institutional memory laundering.

The framework explains why this is so. Static labels are brittle equivalence classes with hidden eps and hidden feature sets. Audit trails are continuity proofs: they show linkage across time, reveal where the chain is strong, and reveal where it breaks. They allow correction without demanding metaphysical certainty.

The framework also predicts characteristic institutional failure modes:

Counterfeit continuity. Systems that permit identity to be asserted without chain evidence create openings for fraud and scapegoating.

Coercive categorization. Systems that collapse people into broad types without recourse will generate predictable injustices: errors that are hard to correct because the category is treated as reality.

Reset theater. Institutions that repeatedly declare new eras (“fresh start,” “new policy,” “reform”) without maintaining continuity of accountability will produce cycles of distrust. The claim “we are different now” requires continuity evidence of constraint change, not a new label.

So the semi-sameness lens suggests a design rule for institutions: treat identity not as a label but as a chain, and treat “same” not as an assertion but as a proof obligation proportional to stakes.

7.3 AI predictions: machine-made semi-sameness and error propagation

AI systems do not merely predict; they classify, match, cluster, and score. That is, they manufacture semi-sameness at scale. They decide which faces are “the same person,” which texts are “the same kind of content,” which users are “the same risk category,” which events are “the same pattern,” which accounts are “the same actor.” Whether or not the system is conscious, it is acting as an equivalence-maker inside social infrastructure.

The semi-sameness framework yields three concrete AI-era predictions.

First, equivalence errors will concentrate in boundary cases where feature sets are incomplete and eps is poorly calibrated. This is not only a technical issue; it is a governance issue. When a system is trained to optimize accuracy on average, it may adopt an eps that is unacceptable for the minority of cases where harm is concentrated. The predictable result is asymmetric harm: certain populations or contexts pay the price for the system’s chosen tolerance.

Second, misbinding will propagate through downstream systems because semi-sameness is compositional. One system’s identity decision becomes another system’s ground truth. A misidentified face can become a legal flag, which becomes a denial of service, which becomes a narrative that “the system says you are this.” The framework predicts that once equivalence classes become institutionalized, they harden. Correction becomes difficult because the equivalence claim is treated as strict identity.

Third, AI will amplify narrative override. Models trained on large corpora can produce fluent, template-driven explanations that force events into familiar categories. This is semi-sameness by storytelling: the system outputs “this is the same as that” because the template is statistically available, not because the continuity evidence is strong. The risk is coherence laundering: a persuasive narrative creates the feeling of sameness without the proof obligations that high-stakes contexts require.

These predictions lead to a sober conclusion: in the AI era, semi-sameness is no longer only a human cognitive necessity. It is an infrastructural force. It can be engineered well, with auditability and recourse, or engineered poorly, with opacity and coercion. The framework’s contribution is to name the mechanism and therefore make its governance legible.

7.4 Design principles: auditability, humility, and repair orientation

The paper's mechanism claim becomes a practical design program once we accept that semi-sameness is unavoidable. The question becomes: how do we build equivalence-makers that are truthful, corrigible, and humane?

Four principles follow directly from the formal discipline introduced earlier.

Principle 1: Auditability of scheme, features, and tolerance. Any system that asserts sameness in high-stakes contexts must be able to disclose, at least in structured form, what feature set F it used, what tolerance eps it applied, and what scheme S governed its identity mapping. This does not require exposing proprietary internals in full detail. It requires making the operational basis contestable. Without this, "same" becomes an unchallengeable decree.

Principle 2: Typed outputs, not monolithic identity claims. Systems should not output a single word like "match" or "same." They should output typed claims: similarity score under features, continuity evidence where relevant, and an explicit scope statement about purpose. This prevents the common failure where an operational equivalence is treated as strict identity simply because the system returned a binary label.

Principle 3: Humility as a system behavior, not a moral slogan. Humility here means calibrated uncertainty and refusal of false closure. In practice, it means thresholds that trigger "insufficient evidence," escalation to human review, and explicit flags where the model's confidence is derived from weak features or distribution shifts. A system that never says "I don't know" will manufacture coerced sameness by necessity.

Principle 4: Repair orientation and recourse. Because equivalence errors are inevitable, ethical systems must be designed around correction. That means appeals processes, reversible decisions where possible, provenance tracking, and mechanisms to update records when misbinding is discovered. Repair orientation treats errors as part of the operating environment rather than as embarrassing exceptions to be hidden.

A fifth principle ties the rest together: proportionality to stakes. The higher the harm potential, the stricter the eps, the richer the feature set, the stronger the continuity proof requirements, and the more robust the recourse mechanisms must be. This principle is the ethical translation of the cognitive point that eps depends on stakes.

Taken together, these principles shift identity-making from rhetoric to disciplined practice. They treat sameness as a tool that must be governed rather than as a metaphysical assumption. And they make a final synthesis visible: consciousness constructs semi-sameness so life can proceed; institutions and AI systems now construct semi-sameness so society can proceed; therefore the

central ethical problem of the AI era is not merely “intelligence” but “equivalence governance” under irreversibility and dignity.

The remaining sections will test the framework against common objections and then close with a concise restatement of what is gained, what is risked, and what work remains.

8. Objections and Replies

The semi-sameness framework invites two predictable reactions. One reaction says it adds nothing: uniqueness is obvious, equivalence is what we already use, and consciousness talk is decorative. The other reaction says it adds too much: it destabilizes the self, smuggles in metaphysical excess, or replaces rigorous explanation with spooky language. This section answers both reactions by tightening the core claims and showing where the framework does real work.

A recurring theme in these replies is the same discipline proposed earlier: treat “same” as a typed claim. Many objections depend on sliding between strict identity and operational sameness, or between descriptive similarity and normative equivalence. Once the types are kept distinct, most objections either dissolve or can be met on their strongest form.

8.1 “This is trivial: everything is unique anyway.”

Objection: Of course everything is unique. Every moment differs. So the thesis is just an elaborate way of saying “things change.”

Reply: The thesis is not merely that things change. It is that the deep human practice of calling things “the same” is not an ontological report about repeated tokens but an engineered equivalence layer needed for agency. That is not trivial because it identifies a mechanism role for consciousness and provides a diagnostic for errors that matter.

A trivial claim does not change what we expect to observe or how we debug failures. This framework does. It predicts specific failure modes when invariants, eps, binding, or continuity chains miscalibrate. It also predicts where institutional and AI systems will create harm by laundering operational equivalence into strict identity. That is a non-trivial explanatory and practical contribution.

The “triviality” objection also ignores the ethical consequence of irreversibility. Saying “everything is unique” in a casual sense is not the same as treating uniqueness as a moral constraint: you cannot redo harm; you cannot erase history; repair must be forward construction. The paper’s thesis is not a slogan; it is a re-anchoring of accountability and forgiveness under a one-pass ontology.

Finally, the framework discloses something that is usually hidden: sameness is manufactured. That is a substantive claim about cognition, language, and institutions. It does not merely state that the world is varied; it explains why "as if sameness" is indispensable and why it is morally dangerous when left undisciplined.

8.2 "Equivalence is all we need; strict identity is idle."

Objection: Physics and science operate on equivalence classes, invariants, and symmetries. They do not need strict identity. Therefore strict identity is a useless metaphysical notion.

Reply: The framework agrees that science primarily uses operational sameness. But it rejects the inference that strict identity is idle, because strict identity plays an essential boundary role in discourse: it prevents category mistakes.

Most real-world confusions occur when people treat operational equivalence as if it were strict identity. They say "the same person," "the same event," or "the same situation," and then draw conclusions that require strict identity: certainty of repeated outcome, entitlement to erase, denial of change, or justification of coercion. Keeping strict identity in the vocabulary makes it possible to say clearly: "that inference is not licensed by equivalence." Strict identity is not needed as a scientific workhorse; it is needed as a conceptual guardrail.

Strict identity also anchors the irreversibility claim. Without it, "repair" can quietly become "rewind." The paper's ethical discipline depends on recognizing that events are strict tokens and therefore cannot be repeated. That does not require strict identity to be common; it requires strict identity to be meaningful as a contrast class.

So the correct stance is not "strict identity is idle," but "strict identity is rarely satisfied across time; operational identity is what we use; and we must not confuse the two." That is exactly the typed-claim discipline the paper advocates.

8.3 "Uniqueness destroys the self."

Objection: If every instant is unique and the self is only a sequence of token-states, then there is no enduring person. This undermines moral responsibility and the meaning of love and commitment.

Reply: This objection assumes that endurance requires strict identity across time. The framework denies that assumption. Endurance requires continuity chains with ownership, not a retained essence.

On this view, the self is not destroyed; it is clarified. The self is the privileged chain of token-states linked by bodily continuity, memory integration, agency control, and consequence inheritance. That chain is exactly what moral responsibility already presupposes in practice. The

framework does not loosen responsibility; it makes it auditable. If the chain is intact, accountability binds. If the chain is disrupted (coercion, incapacity, radical break), accountability is adjusted, as moral practice already does.

Love and commitment also become cleaner, not weaker. Love is revealed as continuity stewardship: holding a person through change without requiring metaphysical stasis, and refusing to reduce them to a single past token. Commitment becomes forward-binding: a promise is not a claim that tomorrow will be the same; it is a claim that the continuity chain will carry an obligation forward.

So uniqueness does not destroy the self. It forbids a simplistic picture of the self as a frozen substance and replaces it with a stronger picture: a temporally extended agent whose identity is maintained by lawful continuity and moral ownership.

8.4 "Structuralism is enough; haecceity is spooky."

Objection: Any appeal to "thisness" sounds like a spooky intrinsic tag. We should adopt structuralism: identity and individuation arise entirely from relations and structures, not from intrinsic haecceities.

Reply: The framework can be expressed entirely in structuralist terms. The minimal schema in Section 6 uses indices, schemes, feature sets, tolerances, and continuity relations. None of this requires a metaphysical thisness-particle.

Haecceity language appears, if at all, only as a phenomenological and ethical pointer: the irreducible "this one" that commands attention and care. But the framework does not use haecceity as an explanatory primitive. The explanatory work is done by continuity chains and equivalence classes under declared schemes.

In fact, structuralism strengthens the framework. It clarifies that "particularity" can be modeled as relational address: a token's position in a causal network, a history, a set of relations that is not duplicated. That captures instantaneous uniqueness without spooky tags.

So the paper does not force a choice between structuralism and haecceity. It recommends a disciplined division of labor: structuralism for ontology and explanation, and a careful allowance that human moral perception needs a term that reminds us not to collapse persons into types. If a reader dislikes haecceity vocabulary, they can drop it without losing any of the paper's mechanism claims or ethical conclusions.

These replies complete the defense of the framework. Semi-sameness is not trivial because it identifies a mechanism role and predicts failures. Strict identity is not idle because it guards against category mistakes. Uniqueness does not destroy the self because selfhood is owned

continuity. Structuralism suffices for the formal scaffold; any haecceity talk is optional and must not do hidden explanatory work.

The remaining work is therefore not to argue endlessly about labels but to apply the framework: to case studies of moral speech, institutional identity systems, and AI-era equivalence engines, and to compile a disciplined Reference section that situates the claims in the relevant literatures.

9. Conclusion: Uniqueness as Ground, Semi-Sameness as Tool

The paper's path has been to start with an austere ground layer—unrepeatable events in a one-pass universe—and then explain how beings like us can live, speak, and bind obligations anyway. The result is a reframing of both metaphysics and ethics. Uniqueness is not a poetic flourish; it is the ontological constraint that prevents "repeat" from being a literal feature of events. Semi-sameness is not an illusion; it is the engineered equivalence layer that makes action and responsibility possible. Consciousness, on this view, is not only a private theater of experience; it is the mechanism that builds and maintains semi-sameness: extracting invariants, assigning tolerances, binding features into objects and agents, and sustaining continuity chains across time.

The conclusion does four things: it restates the thesis tightly, names the main gains, warns against predictable overreach, and identifies the most useful next steps.

9.1 Restate the thesis and mechanism claim

The thesis can be stated as a three-part architecture:

Uniqueness is the ground: token-events and token-states do not recur.

Semi-sameness is the tool: operational "sameness" is an equivalence relation defined by feature choices and tolerances, used for prediction, coordination, and moral accounting.

Consciousness is the builder: consciousness constructs and inhabits semi-sameness by integrating perception, memory, and agency into stable handles and continuity chains.

In the paper's minimal formal terms, the world supplies token-states x(t). Any claim of "same" across times is therefore not strict identity but a typed operational claim. Practical sameness is expressed as x(t1) ~=_(eps, F) x(t2) under declared features and tolerance. Persistence talk becomes continuity talk. The self becomes a privileged continuity class marked by ownership and consequence inheritance. Ethics becomes governance of equivalence-making under stakes: how eps is set, how categories are used, how repair is pursued without erasure.

This restatement is intentionally modest. It does not claim that physics proves a total metaphysics. It does not claim to settle the nature of qualia. It claims that the everyday success of life and the stability of moral practice require an equivalence-making layer, and that consciousness is the most plausible name for the integrative operation that builds that layer in humans.

9.2 What is gained: cleaner metaphysics, ethics, and language

The first gain is metaphysical hygiene. The framework separates what is often conflated:

Token versus type, event versus description, strict identity versus operational identity, persistence versus continuity.

Once those distinctions are explicit, classic confusions become preventable rather than endlessly debatable. Many metaphysical disputes about "identity over time" are revealed as disputes about which continuity criteria matter for which purposes. The framework does not dissolve these disputes by fiat; it makes them legible and therefore tractable.

The second gain is ethical clarity under irreversibility. If events are unrepeatable, ethics becomes forward-directed construction rather than rewind fantasy. Repair becomes a concrete program: acknowledgment, restitution, constraint change, and memory stewardship. Forgiveness becomes a disciplined decision about future equivalence, not a coerced erasure of the past. Accountability becomes continuity-governed rather than essence-governed: obligations bind because the continuity chain binds, not because a metaphysical core remained frozen.

The third gain is linguistic discipline. "Same" stops being a rhetorical blur and becomes a typed claim. Speakers are encouraged—especially in high-stakes contexts—to specify whether they mean strict identity or operational equivalence, what feature set is being compared, what tolerance is being allowed, and what purpose the sameness claim serves. This reduces manipulation and reduces error. It also improves intellectual honesty: we can keep ordinary language while knowing what it is doing.

A fourth gain is modern relevance. In the AI era, semi-sameness is not only a human cognitive artifact; it is an infrastructural force. Machine systems manufacture equivalence classes at scale, and their misbindings propagate through institutions. The framework therefore provides a vocabulary for governance: auditability of features and tolerances, typed outputs instead of monolithic identity labels, humility behaviors that refuse false closure, and repair orientation through recourse and correction.

9.3 What is risked: rhetorical inflation and overreach

The framework has two main risks.

The first is rhetorical inflation of uniqueness. It is easy to turn "everything is unique" into a grand slogan that sounds profound but functions as a conversational solvent: it dissolves distinctions and excuses practical obligations. The paper explicitly refuses that drift. Uniqueness is a ground constraint about token-events; it is not a permission slip to abandon operational sameness where life requires it. The entire framework is built to preserve practical sameness while clarifying its type and scope.

The second risk is overreach about consciousness. Calling consciousness an equivalence-maker can be misread as claiming that consciousness creates reality rather than constructs an interface to reality. The paper does not claim the mind invents the world. It claims the mind constructs stable handles over the world's flux under constraints. If a reader hears "equivalence-maker" as "reality-maker," they have imported a stronger thesis than the paper needs.

A related overreach risk is misusing typed identity talk as a moral escape hatch. Someone might say, "I'm not the same person," intending to void accountability. The framework rejects that misuse. Operational identity can shift in character and capacity, but continuity of agency and consequence inheritance still grounds responsibility unless genuine ruptures can be shown. Typed identity claims are meant to prevent moral laundering, not to enable it.

So the admonition is simple: keep the framework humble. Use it to clarify, not to posture. Use it to demand evidence and repair, not to romanticize flux or dissolve duty.

9.4 Forward work: case studies and a disciplined Reference section

The most valuable next step is case work. The framework should be tested where "same" claims carry real stakes and where equivalence-making systems routinely fail.

Case studies in moral speech. Analyze apologies, confessions, public controversies, and relational repair attempts, identifying where operational equivalence is laundered into erasure, where strictness is demanded selectively, and where truth-labeling discipline would change outcomes.

Case studies in institutional identity. Examine auditing systems, legal identity disputes, credentialing, and administrative continuity claims. Evaluate which designs treat identity as a chain (with provenance and recourse) versus as a static label (with coercion and brittleness).

Case studies in AI-made equivalence. Examine matching, clustering, scoring, and moderation pipelines. Track how a single equivalence decision propagates downstream, and identify which

interventions reduce harm: typed outputs, thresholding to “insufficient evidence,” stronger recourse mechanisms, and logging that supports correction.

Formal refinement. Develop minimal, domain-specific schemes S with declared feature sets F and tolerances eps, and provide templates for how to state operational identity claims in high-stakes settings.

Finally, compile a disciplined Reference section. The references should support the three pillars: (i) event uniqueness and recurrence edge cases, (ii) philosophical accounts of persistence and identity, and (iii) the cognitive, institutional, and AI literature relevant to classification, auditability, and error propagation. The Reference section should be curated for relevance, not maximal. Its job is to ground the paper’s claims without letting citation volume replace conceptual clarity.

The closing synthesis can be stated without flourish: the world does not repeat its events, but life must act as if it can recognize continuity. Consciousness is the mechanism that constructs that recognition as semi-sameness. Ethics is the discipline that governs how we use it. And in the AI era, the question is no longer whether we manufacture sameness, but whether we do so with humility, auditability, and repair.

References

1. Adams, R. M. (1979). Primitive Thisness and Primitive Identity. *The Journal of Philosophy*.
2. Barocas, S., Hardt, M., & Narayanan, A. (2019). *Fairness and Machine Learning: Limitations and Opportunities*. (Book draft widely circulated.)
3. Boltzmann, L. (1872). Weitere Studien über das Wärmegleichgewicht unter Gasmolekülen. *Sitzungsberichte der Kaiserlichen Akademie der Wissenschaften*.
4. Bowker, G. C., & Star, S. L. (1999). *Sorting Things Out: Classification and Its Consequences*. MIT Press. MIT Press Direct
5. Buolamwini, J., & Gebru, T. (2018). Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In *Proceedings of Machine Learning Research* 81. Proceedings of Machine Learning Research
6. Brubaker, R., & Cooper, F. (2000). Beyond "Identity". *Theory and Society*.
7. Cowling, S. (2015). Haecceitism. In *The Stanford Encyclopedia of Philosophy*. Stanford Encyclopedia of Philosophy
8. Cross, R. (2003). Medieval Theories of Haecceity. In *The Stanford Encyclopedia of Philosophy*. Stanford Encyclopedia of Philosophy
9. Dufner, A. (2025). Personal Identity and Ethics. In *The Stanford Encyclopedia of Philosophy*. Stanford Encyclopedia of Philosophy
10. Gallois, A. (2005). Identity Over Time. In *The Stanford Encyclopedia of Philosophy*. Stanford Encyclopedia of Philosophy
11. Grother, P. J., Ngan, M. L., & Hanaoka, K. K. (2019). *Face Recognition Vendor Test (FRVT), Part 3: Demographic Effects*. NIST Interagency Report 8280. NIST Publications
12. Hawking, S. W. (1992). Chronology Protection Conjecture. *Physical Review D*, 46(2), 603–611. APS Physics
13. Hacking, I. (1986). Making Up People. In *Reconstructing Individualism: Autonomy, Individuality, and the Self in Western Thought*.
14. Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux. (Useful for stakes-sensitive thresholds and categorization habits.)
15. Kac, M. (1947). On the Notion of Recurrence in Discrete Stochastic Processes. *Bulletin of the American Mathematical Society*.

16. Kripke, S. (1980). *Naming and Necessity*. Harvard University Press.

17. Lewis, D. (1986). *On the Plurality of Worlds*. Blackwell.

18. Lewis, D. (2001). Reducing Modality to Essence. In *Philosophy and Phenomenological Research*.

19. Locke, J. (1694/1975). *An Essay Concerning Human Understanding* (P. H. Nidditch, Ed.). Oxford University Press. (Classic memory-based identity view.)

20. Olson, E. T. (2002). Personal Identity. In *The Stanford Encyclopedia of Philosophy*. Stanford Encyclopedia of Philosophy

21. Olson, E. T. (1997). *The Human Animal: Personal Identity without Psychology*. Oxford University Press.

22. Olson, E. T. (2007). *What Are We? A Study in Personal Ontology*. Oxford University Press.

23. Parfit, D. (1984). *Reasons and Persons*. Oxford University Press. Stafforini

24. Peirce, C. S. (1878). How to Make Our Ideas Clear. *Popular Science Monthly*. (Operational clarity and meaning via practical consequences.)

25. Poincaré, H. (1890). Sur le problème des trois corps et les équations de la dynamique. *Acta Mathematica*. (Foundational for recurrence/return structure discussions in dynamical systems.)

26. Power, M. (1999; originally 1997). *The Audit Society: Rituals of Verification*. Oxford University Press. OUP Academic

27. Quine, W. V. O. (1960). *Word and Object*. MIT Press. (Reference, indeterminacy, and scheme dependence themes.)

28. Ricoeur, P. (1992). *Oneself as Another*. University of Chicago Press. (Narrative identity and re-identification.)

29. Scotus, J. D. (c. 1300). *Ordinatio* (various editions/translations). (Historical root of haecceity/individuation.)

30. Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and Abstraction in Sociotechnical Systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*.

31. Shoemaker, S., & Swinburne, R. (1984). *Personal Identity*. Blackwell. (Canonical debate framing.)

32. Sider, T. (2001). *Four-Dimensionalism: An Ontology of Persistence and Time*. Oxford University Press.

33. Strawson, P. F. (1959). *Individuals: An Essay in Descriptive Metaphysics*. Methuen. (Particulars, re-identification, and descriptive scheme.)

34. Wiggins, D. (2001). *Sameness and Substance Renewed*. Cambridge University Press.

35. Wiener, N. (1948). *Cybernetics: Or Control and Communication in the Animal and the Machine*. MIT Press. (Continuity, control, and feedback framing relevant to "owned continuity.")